The Structurally Smoothed Graphlet Kernel

Pinar Yanardag
Department of Computer Science
Purdue, West Lafayette IN
ypinar@purdue.edu

S.V. N. Vishwanathan
Dept. of Statistics and Computer Science
Purdue, West Lafayette IN
vishy@stat.purdue.edu

arXiv:1403.0598v1 [cs.LG] 3 Mar 2014

ABSTRACT

A commonly used paradigm for representing graphs is to use a vector that contains normalized frequencies of occurrence of certain motifs or sub-graphs. This vector representation can be used in a variety of applications, such as, for computing similarity between graphs. The graphlet kernel of Shervashidze et al. [32] uses induced sub-graphs of k nodes (christened as graphlets by Przulj [28]) as motifs in the vector representation, and computes the kernel via a dot product between these vectors. One can easily show that this is a valid kernel between graphs. However, such a vector representation suffers from a few drawbacks. As k becomes larger we encounter the sparsity problem; most higher order graphlets will not occur in a given graph. This leads to diagonal dominance, that is, a given graph is similar to itself but not to any other graph in the dataset. On the other hand, since lower order graphlets tend to be more numerous, using lower values of k does not provide enough discrimination ability. We propose a smoothing technique to tackle the above problems. Our method is based on a novel extension of Kneser-Ney and Pitman-Yor smoothing techniques from natural language processing to graphs. We use the relationships between lower order and higher order graphlets in order to derive our method. Consequently, our smoothing algorithm not only respects the dependency between sub-graphs but also tackles the diagonal dominance problem by distributing the probability mass across graphlets. In our experiments, the smoothed graphlet kernel outperforms graph kernels based on raw frequency counts.

1. INTRODUCTION

In this paper, we are interested in comparing graphs by computing a kernel between graphs [39]. Graph kernels are popular because many datasets from diverse domains such as bio-informatics [3, 30], chemo-informatics [2], and web data mining [41] naturally can be represented as graphs. Almost all graph kernels (implicitly or explicitly) represent a graph as a (normalized or un-normalized) vector which contains the frequency of occurrence of motifs or sub-graphs[1]. The key idea here is that well chosen motifs can capture the semantics of the graph structure while being computationally tractable. For instance, counting walks in a graph leads to the random walk graph kernel of Borgwardt et al. [3] (see Vishwanathan et al. [39] for an efficient algorithm for computing this kernel). Other popular motifs include subtrees [31], shortest paths [5], and cycles [16]. Of particular interest to us are the graphlet kernels of Shervashidze et al. [32]. The motif used in this kernel is the set of unique sub-graphs of size k, which were christened as graphlets by Przulj [28].

[1]The kernels proposed by [21] are a notable exception.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
KDD'14, August 24–27, 2014, New York City, New York USA.
Copyright 2014 ACM 978-1-4503-1462-6 /12/08 ...$15.00.

Observation 1: Computing meaningful graphlet kernels that have high discriminative ability requires a careful selection of k. If k is small, then the number of unique graphlets is small (i.e., the length of the feature vector is small). See Figure 1. Consequently the feature vector does not provide meaningful discrimination between two graphs. On the other hand, if k is large then a) the set of unique graphlets grows exponentially (i.e., the feature vector is very high dimensional) but b) only a small number of unique graphlets will be observed in a given graph (i.e., the feature vector is very sparse). Moreover, the probability that two graphs will contain a given large sub-graph is very small. Consequently, a graph is similar to itself but not to any other graph in the training data. This is well known as the diagonal dominance problem in the machine learning community [17], and the resulting kernel matrix is close to the identity matrix. In other words, the graphs are orthogonal to each other in the feature space. However, it is desirable to use large values of k in order to gain better discriminative ability. One way to circumvent the diagonal dominance problem is to view the normalized graphlet-frequency vector as estimating a multinomial distribution, and use smoothing.

Observation 2: The normalized graphlet-frequency vector exhibits power-law behavior, especially for large values of k. In other words, a few popular graphlets occur very frequently while a vast majority of graphlets will occur very rarely. Put another way, a few graphlets dominate the distribution. To see this, we randomly sampled a graph from six benchmark datasets (details of the datasets can be found in Section 6), and exhaustively computed occurrences of all graphlets of size $k = 8$ and plotted the resulting histogram on a log-log scale in Figure 2. As can be seen, the frequencies are approximately linear in the log-log scale which indicates power law behavior. Therefore, any smoothing technique that we use on the normalized graphlet-frequency vector must respect this power-law behaviour.

Observation 3: The space of graphlets is structured. What we mean by this is that graphlets of different sizes are related

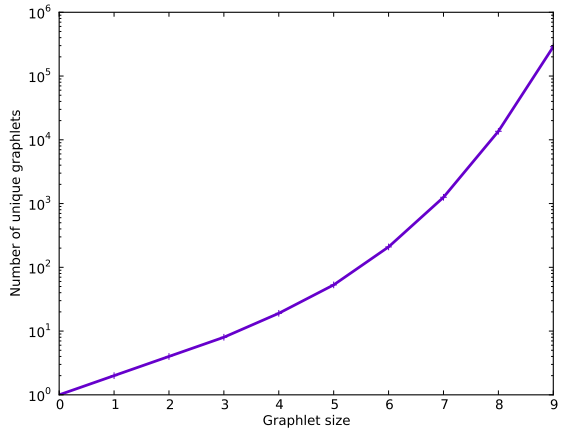

Figure 1: Number of unique graphlets increase exponentially with graphlet size k.

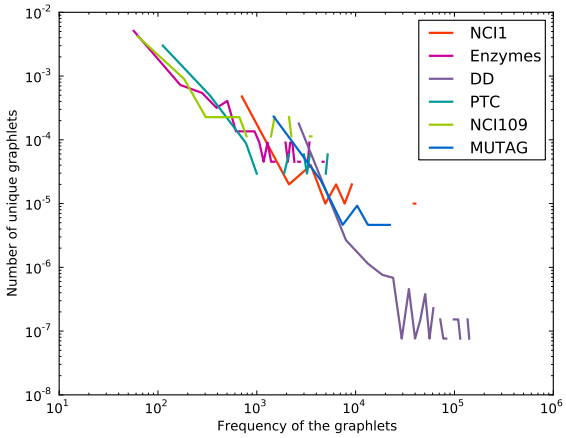

Figure 2: We randomly select a graph from six benchmark graph datasets and exhaustively searched for all graphlets of size $k = 8$. The histogram is plotted in log-log scale in order to demonstrate the power-law behaviour.

to each other. While many such relationships can be derived, we will work with perhaps the simplest one which is depicted in Figure 4. Here, we construct a directed acyclic graph (DAG) with the following property: a node at depth k denotes a graphlet of size k. Given a graphlet g of size k and other graphlet g' of size $k+1$ we add an edge from g to g' if, and only if, g can be obtained from g' by deleting a node of g'. This shows that graphlets of size k have a strong relationship to graphlets of size $k+1$ and one must respect this relationship when deriving a smoothing technique.

Our contributions in this paper are as follows. First, we propose a new smoothing technique for graphlets which is inspired by Kneser-Ney smoothing [20] used for language models in natural language processing. Our model satisfies the desiderata that we outlined above, that is, it respects the power law behavior of the counts and yet takes into account the structure of the space of graphlets. Second, we provide a novel Bayesian version of our model that is extended from the Hierarchical Pitman-Yor process of Teh [35]. Unlike the traditional Hierarchical Pitman-Yor Process (HPYP) where the base distribution is given by another Pitman-Yor Process (PYP), in our case it is given by a transformation of a PYP that is guided by the structure of the space. Third, we perform experiments to validate and understand how smoothing affects the performance of graphlet kernels.

The structure of the paper is as follows. In Section 2, we discuss background on graphlet kernels and smoothing techniques. In Section 3, we introduce our Kneser-Ney-inspired smoothing technique. In Section 4, we propose an alternate Bayesian version of our model. Related work is discussed in Section 5. In Section 6, we perform experiments and discuss our findings, and we conclude the paper with Section 7.

2. BACKGROUND

2.1 Notation

A *graph* is a pair $G = (V, E)$ where $V = \{v_1, v_2, \ldots, v_{|V|}\}$ is an ordered set of *vertices* or *nodes* and $E \subseteq V \times V$ is a set of *edges*. Given $G = (V, E)$ and $H = (V_H, E_H)$, H is a *sub-graph* of G iff there is an injective mapping $\alpha : V_H \to V$ such that $(v, w) \in E_H$ iff $(\alpha(v), \alpha(w)) \in E$. Two graphs $G = (V, E)$ and $G' = (V', E')$ are *isomorphic* if there exists a bijective mapping $g : V \to V'$ such that $(v_i, v_j) \in E$ iff $(g(v_i), g(v_j)) \in E'$. *Graphlets* are small, connected, non-isomorphic sub-graphs of a large network. They were introduced by Przulj [28] to design a new measure of local structural similarity between biological networks. Graphlets up to size five are shown in Figure 3.

2.2 The Graphlet kernel

Let $\mathcal{G}_k = \{g_1, g_2, \ldots, g_{n_k}\}$ be the set of size-k graphlets where n_k denotes the number of unique graphlets of size k. Given a graph G, we define f_G as a normalized vector of length n_k whose i-th component corresponds to the frequency of occurrence of g_i in G:

$$f_G = \left(\frac{c_1}{\sum_j^{n_k} c_j}, \cdots, \frac{c_{n_k}}{\sum_j^{n_k} c_j}\right)^T. \quad (1)$$

Here c_i denotes number of times g_i occurs as a sub-graph of G. Given two graphs G and G', the graphlet kernel k_g is defined as:

$$k_g(G, G') := f_G^\top f_{G'}, \quad (2)$$

which is simply the dot product between the normalized graphlet-frequency vectors.

2.3 Smoothing multinomial distributions

In this section we will briefly review smoothing techniques for multinomial distributions and show that graphlet kernels are indeed based on estimating a multinomial. Suppose we observe a sequence $e_1, e_2, \ldots, e_n$ containing n discrete events drawn from a ground set of size M and we would like to estimate the probability $P(e_i)$ of observing each event e_i. Maximum likelihood estimation based on sequence counts obtained from the observations provides a way to compute $P(e_i)$:

$$P_{MLE}(e_i) = \frac{c_i}{\sum_j c_j}, \quad (3)$$

where c_i denotes the number of times the event e_i appears in the observed sequence and $\sum_j c_j$ denotes the total number of observed events. Therefore, one can easily see that the representation used in graphlet kernels in Section 2.2 is actually an MLE estimate on the observed sequences of graphlets.

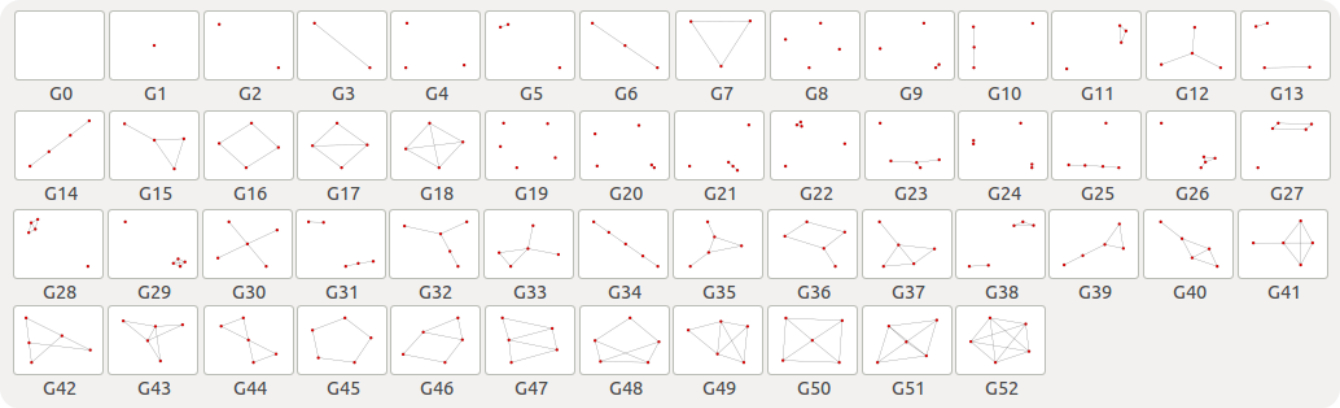

Figure 3: Connected, non-isomorphic induced sub-graphs of size $k \leq 5$. Plots are generated with NetworkX library [14].

However, MLE estimates of the multinomial distribution are *spiky*, that is, they assign zero probability to events that did not occur in the observed sequence. What this means is that an event with low probability is often estimated to have zero probability mass. This issue occurs in a number of different domains and therefore, unsurprisingly, has received significant research attention [42]; smoothing methods are typically used to address this problem. The general idea behind smoothing is to discount the probabilities of the observed events and to assign extra probability mass to unobserved events.

Laplace smoothing is the simplest and one of the oldest smoothing methods, where only a fixed count of 1 is added to every event. This results in the estimate

$$P_{Laplace}(e_i) = \frac{\sum_j c_j}{\sum_j c_j + M} P_{MLE} + \frac{M}{\sum_j c_j + M} \frac{1}{M} \quad (4)$$

or equivalently,

$$P_{Laplace}(e_i) = \lambda P_{MLE}(e_i) + (1 - \lambda)\frac{1}{M}, \quad (5)$$

where λ is a normalization factor which ensures that the distributions sum to one. The intuition behind Laplace smoothing is basically to interpolate a uniform distribution with the MLE distribution. Although Laplace smoothing resolves the zero-count problem, it does not produce a power law distribution which is a desirable feature in real-life models. Therefore, researchers have worked on finding smoothing techniques that respect power law behavior. The key idea behind these methods is to redistribute the probability mass using a so-called *fallback* model, where the fallback model is also recursively estimated.

Kneser-Ney smoothing is a fallback based smoothing method which has been identified as the state-of-the-art smoothing in natural language processing by several studies [7]. Kneser-Ney smoothing computes the probability of an event by using the raw counts that are discounted by using a fixed mass. Then, the discounted mass is re-added equally to all event probabilities by using a base distribution:

$$P_{KN}(e_i) = \frac{max\{c_i - d, 0\}}{\sum_j c_j} + \sum_{j=1}^{n} |\{e_j : c_j > d\}| \frac{d}{\sum_j c_j} P_0(e_i), \quad (6)$$

where $d \geq 0$ is the discounting parameter, $P_0(\cdot)$ is the base distribution and $P_0(e_i)$ denotes the probability mass the base distribution assigned to event e_i. The quantity $\sum_{j=1}^{n} |\{e_j : c_j > d\}|$ is a normalization factor to ensure that the distribution sums to 1 and simply denotes the number of events the discount is applied. When discount parameter $d = 0$, we recover MLE estimation since no mass is taken away from any event. When d is very large, then we recover the base distribution on the events since we discount all the available mass. One should interpolate between these two extremes in order to get a reasonable smoothed estimate. In order to propose a new Kneser-Ney-based smoothing framework, one needs to specify the discount parameter d and the base distribution $P_0(\cdot)$. In the next section we will show how one can derive a meaningful base distribution for graphlets.

3. DEFINING A BASE DISTRIBUTION

The space of graphlets has an inherent structure. One can construct a directed acyclic graph (DAG) in order to show how different graphlets are related with each other. A node at depth k denotes a graphlet of size k. When it is clear from context, we will use the node of the DAG and a graphlet interchangeably. Given a graphlet g_i of size k and another graphlet g_j of size $k+1$ we add an edge from g_i to g_j if, and only if, g_i can be obtained from g_j by deleting a node of g_j. We first discuss how to construct this DAG and then discuss how we can use this DAG to define a base distribution.

The first step towards constructing our DAG is to obtain all unique graphlet types of size k. Therefore, we first exhaustively generate all possible graphs of size k (this involves a one time $O(2^k)$ effort), and use Nauty [25] to obtain their canonically-labelled isomorphic representations. In order to obtain the edges of the DAG, we take a node at depth $k+1$, which denotes a canonically-labeled isomorphic graphlet and delete a node to obtain a size k graphlet. We use Nauty to canonically-label the size k graph, which in turn allows us to link to a node at depth k. By deleting each node of the $k+1$ sized graphlet we can therefore obtain $k+1$ possible links. We repeat this for all nodes at level $k+1$ before proceeding to level k. Figure 4 shows the constructed DAG for size $k = 5$ graphlets. Since all descendants of a given graphlet at level k are at level $k+1$, a topological ordering of the vertices is possible, and hence it is easy to see that the resulting graph is a DAG.

Now, let us define the edge weight between an arbitrary graphlet g_j of size $k+1$ and its parent g_i of size k. Let s_{ij} denote the number of times g_i occurs as a sub-graph of g_j and $\mathcal{C}_{g_i}$ denote all the children of graphlet g_i in the DAG. Then, we define the edge

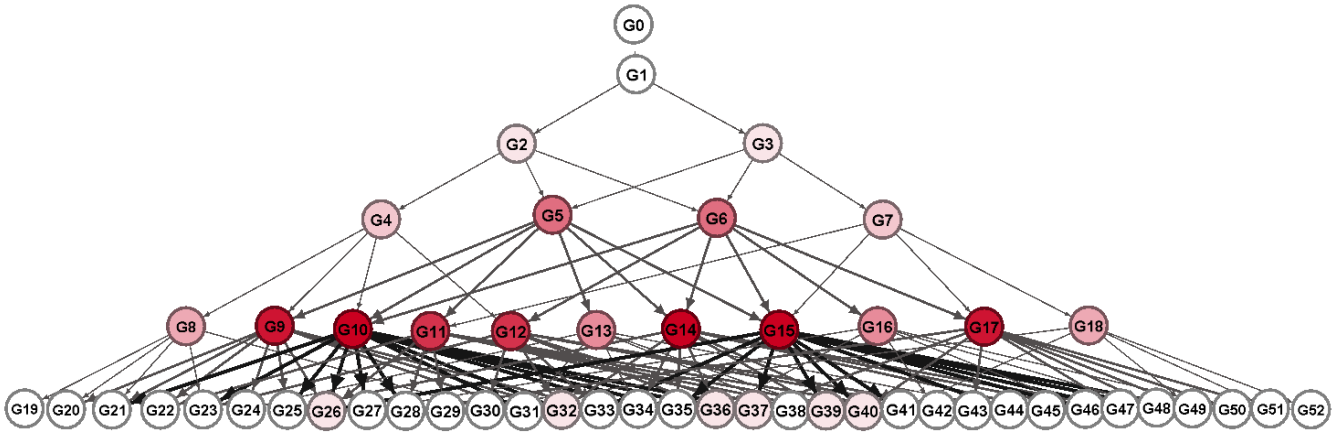

Figure 4: Topologically sorted graphlet DAG where each node are within one delete-distance away. Nodes are colored by average degree. Image is generated by Gephi [1].

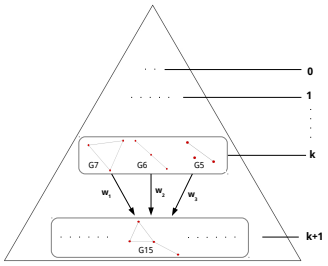

Figure 5: Graphlet g_{15} get the probability mass to its parents g_7, g_6, g_5 according to the weights w_1, w_2, w_3 respectively.

weight between graphlet g_i and g_j as

$$w_{ij} = \frac{s_{ij}}{\sum_{g_{j'} \in \mathcal{C}_{g_i}} s_{ij'}} \tag{7}$$

Next we show how the DAG can be used to define a base distribution. Suppose we have a distribution over graphlets of size k. Then we can transform it into a distribution over size $k+1$ graphlets in a recursive way by exploiting edge connections in the DAG as follows:

$$P_0(g_j) = \sum_{g_i \in Pa(g_j)} w_{ij} P_0(g_i). \tag{8}$$

Here g_j denotes a graphlet of size $k+1$ and $Pa(g_j)$ denotes the parents of graphlet g_j in the DAG.

LEMMA 1. *Given the set of child nodes of a graphlet g_i, if the edge weights on the DAG are all non-negative and satisfy*

$$\sum_{g_j \in C(g_i)} w_{ij} = 1, \tag{9}$$

then (8) *defines a valid probability distribution.*

PROOF. For clarity, we introduce a few notations to facilitate the proof. Assume that there are in total J child graphlets, which we denote by $g_1, g_2, \cdots, g_J$. Further assume that there are in total I parent graphlets, which we denote by $g_1, \cdots, g_I$. Let I_j denote the number of parents graphlet g_j has, i.e. $I_j = |Pa(g_j)|$. Clearly, we have that $\sum_{j=1}^{J} I_j = I$. Thus, we have

$$\begin{aligned}
\sum_{j=1}^{J} P_0(g_j) &= \sum_{j=1}^{J} [\sum_{i=1}^{I_j} w_{ij} P_0(g_i)] \\
&= \sum_{i=1}^{I} [\sum_{\{j: g_j \in C(g_i)\}} w_{ij} P_0(g_i)] \\
&= \sum_{i=1}^{I} P_0(g_i) [\sum_{\{j: g_j \in C(g_i)\}} w_{ij}] \\
&= \sum_{i=1}^{I} P_0(g_i) = 1,
\end{aligned}$$

which completes the proof. □

The base distribution we defined above respects the structural space of the graphlets. Pretend for a moment that we are given only the frequencies of occurrences of size k graphlets and are asked to infer the probability of occurrences of size $k+1$ graphlets. Without any additional information, one can infer the distribution as follows: each parent graphlet casts a vote for every child graphlet based on how many times the parent occurs in the child. The votes of all parents are accumulated and this provides a distribution over size $k+1$ graphlets. In other words, a natural way to infer the distribution at level $k+1$ is to use how likely we are to see its sub-graphs. Figure 5 illustrates the relationship between a graphlet of size $k+1$ and its parent graphlets of size k. Here, edge weights denote how many times each parent occurs as a sub-graph of g_{15}. In the case that we do not observe graphlet g_{15}, it still gets probability mass proportional to the edge weight from its parents g_7, g_6, g_5, thus overcoming the sparsity problem of unseen data. Our model combines this base distribution with the observed real data to generates the final distribution.

The way how the discounted mass is distributed is controlled by the edge weights between two graphlets. In Equation 7 we defined edge weights according to the number of times a parent node occurs in its children. However, one can explore different weighting schemes between the nodes on the DAG based on domain knowl-

edge. For example, in the case of structured graphs such as social networks, one might benefit from weighting the edges according to the PageRank [26] score of the nodes. Similarly, other link analysis algorithms such as Hubs or Authority given by HITS algorithm [19] can be used in order to exploit the domain knowledge.

3.1 Kneser-Ney Smoothing with a structural distribution

We now have all the components needed to explain our Structural Kneser-Ney (SKN) framework. Given an arbitrary graphlet g_j of size $k+1$, we estimate the probability of observing that graphlet as follows:

$$P_{SKN}(g_j) = \frac{max(c_j - d, 0)}{\sum_{g_{j'} \in \mathcal{G}_{k+1}} c_{j'}} + \frac{d}{\sum_{g_{j'} \in \mathcal{G}_{k+1}} c_{j'}} \sum_{g_{j'} \in \mathcal{G}_{k+1}} |\{g_{j'} : c_{j'} > d\}| \sum_{i \in \mathcal{P}_{g_j}} P_0(g_i) \frac{w_{ij}}{\sum_{g_{j'} \in \mathcal{C}_{g_i}} w_{ij'}} \quad (10)$$

As can be seen from the equation, we first discount the count of all graphlets by d, and then redistribute this mass to all other graphlets. The amount of mass a graphlet receives is controlled by the base distribution. In order to automatically tune the discount parameter d, we use the Pitman-Yor process (a Bayesian approximation of Kneser-Ney) in the next section.

4. PITMAN-YOR PROCESS

We will only give a very high level overview of a Pitman-Yor process and refer the reader to the excellent papers by Teh [35] and [12] for more details. A Pitman-Yor process P on a ground set $\mathcal{G}_{k+1}$ of size-$(k+1)$ graphlets is defined via

$$P_{k+1} \sim PY(d_{k+1}, \theta_{k+1}, P_k), \quad (11)$$

where d_{k+1} is a discount parameter $0 \leq d_{k+1} < 1$, $\theta > -d_{k+1}$ is a strength parameter, and P_k is a base distribution. The most intuitive way to understand draws from the Pitman-Yor process is via the Chinese restaurant process (also see Figure 6). Consider a restaurant with an infinite number of tables. Customers enter the restaurant one by one. The first customer sits at the first table, and is seated at the first table. Since this table is occupied for the first time, a graphlet is assigned to it by drawing a sample from the base distribution. The label of the first table is the first graphlet drawn from the Pitman-Yor process. Subsequent customers when they enter the restaurant decide to sit at an already occupied table with probability proportional to $c_i - d_{k+1}$, where c_i represents the number of customers already sitting at table i. If they sit at an already occupied table, then the label of that table denotes the next graphlet drawn from the Pitman-Yor process. On the other hand, with probability $\theta_{k+1} + d_{k+1}t$, where t is the current number of occupied tables, a new customer might decide to occupy a new table. In this case, the base distribution is invoked to label this table with a graphlet. Intuitively the reason this process generates power-law behavior is because popular graphlets which are served on tables with a large number of customers have a higher probability of attracting new customers and hence being generated again. This self reinforcing property produces power law behavior.

In a hierarchical Pitman-Yor process, the base distribution P_k is recursively defined via a Pitman-Yor process $P_k \sim PY(d_k, \theta_k, P_{k-1})$. In order to label a table, we need a draw from P_k, which is obtained by inserting a customer into the corresponding restaurant.

In our case P_{k+1} is defined over $\mathcal{G}_{k+1}$ of size n_{k+1} while P_k is defined over $\mathcal{G}_k$ of size $n_k \leq n_{k+1}$. Therefore, like we did in the case of Kneser-Ney smoothing we will use the DAG and Equation (8) to define a base distribution. This changes the Chinese Restaurant process as follows: When we need to label a table, we will first draw a size-k graphlet $g_i \sim P_k$ by inserting a customer into the corresponding restaurant. Given g_i, we will draw a size-$(k+1)$ graphlet g_j proportional to w_{ij}, where w_{ij} is obtained from the DAG. Deletion of a customer is handled similarly. Detailed pseudo-code can be found in Algorithms 1 and 2.

Algorithm 1 Insert a Customer

Input: $d_{k+1}, \theta_{k+1}, P_k$
  $t \leftarrow 0$ // Occupied tables
  $c \leftarrow ()$ // Counts of customers
  $l \leftarrow ()$ // Labels of tables
  if $t = 0$ **then**
    $t \leftarrow 1$
    append 1 to c
    draw graphlet $g_i \sim P_k$ // Insert customer in parent
    draw $g_j \sim w_{ij}$
    append g_j to l
    return g_j
  else
    with probability $\propto \max(0, c_j - d)$
    $c_j \leftarrow c_j + 1$
    return l_j
    with probability proportional to $\theta + dt$
    $t \leftarrow t + 1$
    append 1 to c
    draw graphlet $g_i \sim P_k$ // Insert customer in parent
    draw $g_j \sim w_{ij}$
    append g_j to l
    return g_j
  end if

Algorithm 2 Delete a Customer

Input: d, θ, P_0, C, L, t
  with probability $\propto c_l$
  $c_l \leftarrow c_l - 1$
  $g_j \leftarrow l_j$
  if $c_l = 0$ **then**
    $P_k \propto 1/w_{ij}$
    delete c_l from c
    delete l_j from l
    $t \leftarrow t - 1$
  end if
  return g

5. RELATED WORK

The problem of estimating multinomial distributions is a classic problem. In natural language processing they occur in the following context: suppose we are given a sequence of words $w_1, \ldots, w_k$ and one is interested in asking what is the probability of observing word w next. Estimating this probability lies at the heart of language models, and many sophisticated smoothing techniques have been proposed. This is a classic multinomial estimation problem that suffers from sparsity since the event space is unbounded. Moreover, natural language exhibits power law behavior since the distribution tends to be dominated by a small number of frequently occurring words. In extensive empirical evaluation it has been found the Kneser-Ney smoothing is very effective for language

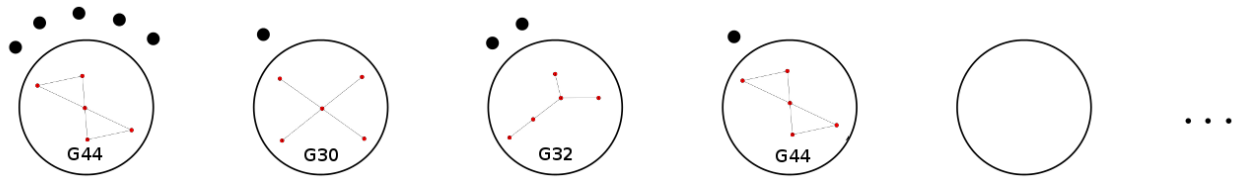

Figure 6: An illustration of table assignment, adapted from Goldwater et al. [13]. In this example, labels at the tables are given by $(l_1, \ldots, l_4) = (G_{44}, G_{30}, G_{32}, G_{44})$. Black dots indicate the number of occurrences of each label in 10 draws from the Pitman-Yor process.

models [7], [24]. Here, the base distribution is constructed using smaller context of $k-1$ words which naturally leads to a denser distribution. Even though language models and graphlets have some similarities, there is a significant fundamental difference between the two. In language models, one can derive the base distribution using a smaller context. However, in the case of graphlets there is no equivalent concept of a fallback model. Therefore, we need to derive the base distribution by using *smaller size* graphlets. However, this leads to a problem since the distribution is now defined on a smaller space. Therefore, we need to apply a transformation by using the DAG in order to convert the distribution back into to the original space. Goldwater et al. [13] and Teh [35] independently showed that Kneser-Ney can be explained in a Bayesian setting by using the Pitman-Yor Process (PYP) [27]. In the Bayesian interpretation, a hierarchical PYP where the Pitman-Yor prior comes from another PYP is used. Similar to Kneser-Ney, this interpretation is not directly applicable to our model since the previous PYP has a different space, thus we need to apply a transformation.

Graph kernels can be considered as special cases of convolutional kernels proposed by Haussler [15]. In general, graph kernels can be categorized into three classes: graph kernels based on walks and paths [11], [18],[5], graph kernels based on limited-size subgraphs [16], [31], [34] and graph kernels based on subtree patterns [29]. Shervashidze et al. [33] performs a relaxation on the vertices and exploit labeling information embedded in the graphs to derive their so-called Weisfeiler-Lehman kernels. However, their kernel is applicable only to labeled graphs. The sparsity problem of graphlet kernels has been addressed before. Hash kernels proposed by Shi et al. [34] addresses the sparsity problem by applying a sparse projection into a lower dimensional space. The idea here is that many higher order graphlets will "collide" and therefore be mapped to the same lower dimensional representation, thus avoiding the diagonal dominance problem. Unfortunately, we find that in our experiments the hash kernel is very sensitive to the hash value used for embedding and rarely performed well as compared to the MLE estimate.

6. EXPERIMENTS

To compare the efficacy of our approach, we compare our Kneser-Ney and Pitman-Yor smoothed kernels with state-of-the-art graph kernels namel the graphlet kernel [32], the hash kernel [34], the random walk kernel [11], [18], [38], and the shortest path kernel [5]. For random walk kernel, we uniformly set the decay factor $\lambda = 10^{-4}$, for shortest path we used the delta kernel to compare the shortest-path distances, and for the hash kernel we used a prime number of 11291. We adopted Markov chain Monte Carlo sampling based inference scheme for the hierarchical Pitman-Yor language model from [35] and modified the open source implementation of HPYP from `https://github.com/redpony/cpyp`.

Table 2: Properties of the datasets

Dataset	Size	Classes	Avg Nodes	Avg Edges
MUTAG	188	2 (125 vs 63)	17.9	39.5
PTC	344	2 (152 vs 192)	25.5	51.9
Enzyme	600	6 (100 each)	32.6	124.2
DD	1178	2 (691 vs 487)	284.3	1431.3
NCI1	4110	2 (2057 vs 2053)	29.8	64.6
NCI109	4127	2 (2079 vs 2048)	29.6	62.2

Due to lack of space we will only present a subset of our experimental results. Full results including the source code and experimental scripts will be made available for download from `http://cs.purdue.edu/~ypinar/kdd`.

Datasets

In order to test the efficacy of our model, we applied smoothing to real-world benchmark datasets, namely MUTAG, PTC, NCI1, NCI109, ENZYMES and DD. MUTAG [9] is a binary data set of 188 mutagenic aromatic and heteroaromatic nitro compounds, labeled whether they have mutagenicity in *Salmonella typhimurium*. The Predictive Toxicology Challenge (PTC) [36] dataset is a chemical compound dataset that reports the carcinogenicity for male and female rats. NCI1 and NCI109 [40], (http://pubchem.ncbi.nlm.nih.gov) datasets, made publicly available by the National Cancer Institute (NCI), are two subsets of balanced data sets of chemical compounds screened for ability to suppress or inhibit the growth of a panel of human tumor cell lines. Enzymes is a data set of protein tertiary structures obtained from [4]. DD [10] is a data set of protein structures where each protein is represented by a graph and nodes are amino acids that are connected by an edge if they are less than 6 Angstroms apart. Table 2 shows summary statistics for these datasets. Note that we did not use edge or node labels in our experiments.

Experimental Setting

All data sets we work with consist of sparse graphs. However, counting all graphlets of size k for a graph with n nodes requires $O(n^k)$ effort which is intractable even for moderate values of k. Therefore, we use random sampling, as advocated by [32], in order to obtain an empirical distribution of graphlet counts that is close to the actual distribution of graphlets in the graph. For each value of for each $k \in \{2, \ldots, 8\}$ we randomly sampled 10,000 sub-graphs, and used Nauty [25] to get canonically-labeled isomorphic representations which are then used to construct the frequency representation.

Table 1: Classification accuracy with standard deviation on benchmark data sets. RW: random walk kernel, SP: shortest path kernel, GK: graphlet kernel, HK: hash kernel, KN: Kneser-Ney smoothing, PYP: Pitman Yor smoothing, >24h: computation did not finish within 24 hours. GK, KN and PYP results are reported for size 5 graphlets.

Dataset	RW	SP	GK	HK	KN	PYP
MUTAG	83.51	78.72	80.34±3.0	80.34±3.0	**82.98** ±5.4	81.94±2.9
PTC	51.16	50	57.26±4.6	57.26±4.6	**59.87**±4.6	56.36±4.3
DD	> 24h	> 24h	72.74±1.9	72.74±1.9	**74.95**±2.3	73.51±2.8
ENZYMES	18.5	21.66	19.50 ±4.4	19.50 ±4.4	**25.66** ±1.6	23.33 ±3.4
NCI1	44.84	63.65	56.56±4.5	56.56±4.5	**62.40** ±1.5	61.60±2.6
NCI109	59.80	62.44	62.00±4.0	62.00±4.0	**62.15** ±1.8	58.32±4.6

Table 3: Graphlet Kernel with exhaustive sampling vs. Smoothed Kernel with 10,000 samples

Dataset	GK	KN	PYP
MUTAG	80.85	**82.98** ±5.4	81.94±2.9
PTC	54.94	**59.87**±4.6	56.36±4.3
DD	> 48h	**74.95**±2.3	73.51±2.8
ENZYMES	> 48h	**25.66** ±1.6	23.33 ±3.4
NCI1	62.36	**62.40** ±1.5	61.60±2.6
NCI109	**62.53**	62.15 ±1.8	58.32±4.6

We performed 5-fold cross-validation with C-Support Vector Machine Classification using LibSVM [6], using 4 folds for training and 1 for testing. We used a linear kernel. Since we are interested in understanding the difference in performance between smoothed and un-smoothed kernels we did not tune the value of C; it was simply set to 1. In order to tune the discount parameter for Kneser-Ney based smoothed kernel, we tried different parameters vary from 0.01 to 10,000 and report results for the best one.

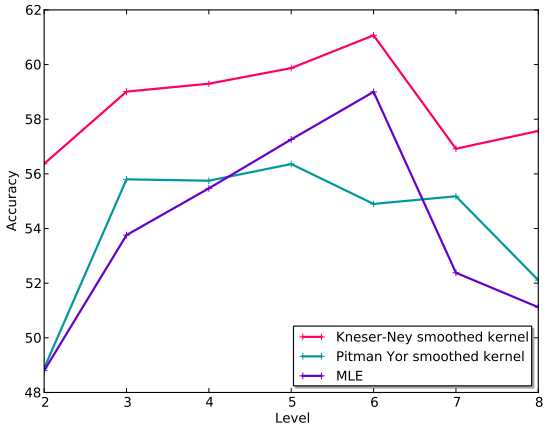

Figure 7: Accuracy vs. graphlet size for PTC dataset with MLE, PYP and Kneser-Ney smoothed kernels

Effect of discounting parameter on performance

First, we investigate the effect of the discounting parameter on the classification performance. Since the trends are similar across different datasets, we pick PTC as a representative dataset to report results. Figure 8 shows the classification accuracy on the PTC dataset with different discounting parameters for Kneser-Ney smoothing. As expected, applying very large discounts decreases the performance because the distribution (10) degrades to the base distribution. On the other hand, applying a very small discount also decreases the accuracy since the distribution degrades to the MLE estimate. From our experiments, we observe that the best performance in all datasets is achieved by using an intermediate discount value between these two extremes. However, the specific discount value is data dependent.

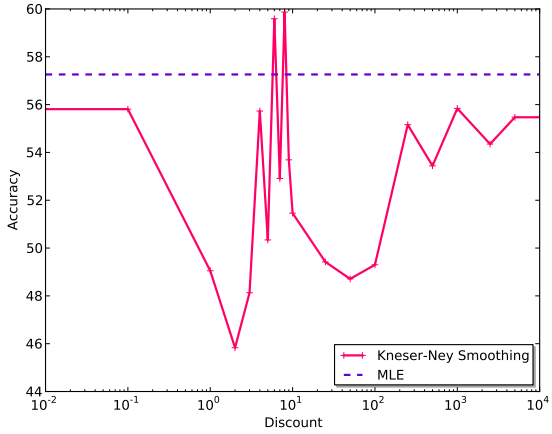

Figure 8: Accuracy vs. Discounting for PTC dataset, k=5 graphlets

Effect of graphlet size on performance

Next we investigate how the value of k affects performance. Again, we show results for a representative dataset namely PTC. Figure 7 shows the classification accuracy on the PTC dataset as a function of graphlet size k for MLE (the graphlet kernel), Pitman-Yor smoothed kernel and Kneser-Ney smoothed kernel. From the figure, we can see that small graphlet sizes such as $k = 2, 3$ do not perform well and are not informative since the number of unique graphlets is very small (see Table 1). On the other hand, MLE does

not perform well for large graphlet sizes such as $k = 7, 8$ because of the diagonal dominance problem. On the other hand, smoothed kernels in general obtain a balance between these two extreme situations and tend to yield better performance. Pitman-Yor smoothed kernel tends to achieve a better performance than MLE, but doesn't perform as good as Kneser-Ney. This is expected since we didn't tune the hyperparameters for Pitman-Yor process. Teh [35] shows that Pitman-Yor yields a better performance if one tune the hyperparameters. Therefore, our Pitman-Yor kernel is open to improvement.

Comparison with related work

We compare the proposed smoothed kernel with graphlet kernel and hash kernel on the benchmark data sets in Table 1. The results for the Shortest Path and Random Walk graph kernels are included mainly to show what is the state-of-the-art using other representations. We fixed the $k = 5$ which is observed to be the best value for the MLE based graphlet kernel on most datasets. We randomly sample 10,000 graphlets from each graph and feed the same frequency vectors to graphlet kernel (GK), hash kernel (HK), Kneser-Ney smoothed kernel (KN), and Pitman-Yor smoothed kernel (PYP). Therefore, the differences in performance that we observe are solely due to the transformation of the frequency vectors that these kernels perform. We performed an unpaired t-test and use bold numbers to indicate that the results were statistically significant at $p < 0.0001$.

We can see that KN kernel outperforms MLE and hash kernels on all of the benchmark data sets. The accuracy of the PYP kernel is usually lower than that of the KN kernel. We conjecture that this is because the Pitman-Yor process is sensitive to the hyperparameters and we do not carefully tune the hyper-parameters in our experiments. On PTC, DD and Enzymes datasets, the KN kernel reached the highest accuracy. On MUTAG, NCI1 and NCI109 datasets, KN kernels also yield good results and got comparable classification accuracies to shortest path and random walk kernels. For the DD dataset, shortest path and random walk kernels were not able to finish in 24 hours, due to the fact that this dataset has a large maximum degree.

To summarize, smoothed kernels turns out to be competitive in terms of classification accuracy on all datasets and are also applicable to very large graphs.

Effect of exhaustive sampling on performance

Next, we investigate whether any of the difference in performance can be attributed to sampling a small number of graphlets. In other words, we ask do the results summarily change if we performed exhaustive sampling instead of using 10,000 samples. We give MLE an unfair advantage by performing exhaustive sampling on MUTAG, PTC, NCI and NCI109 datasets for $k = 5$ by using a distributed memory implementation. Table 4 shows mean, median and standard deviations of number of samples in bruteforce sampled datasets for $k = 5$. Here, we can see that the original frequencies of the graphlets are quite high in most of the datasets. Even though our algorithm only uses 10,000 samples, it outperforms the graphlet kernel with exhaustive sampling on MUTAG, PTC and NCI1 and achieves competitive performance on NCI109 dataset. Results are summarized in Table 3.

Even though distribution with a small number of samples is close to the original distribution in the L_1 sense, bruteforce sampling reveals that the true underlying distribution of the datasets contains a larger number of unique graphlets comparing to random sampling. Since the graphlet kernel uses a MLE estimate its performance degrades. On the other hand, our smoothing technique uses structural information to redistribute the mass and hence is able to outperform MLE even with a small number of samples.

Table 4: Mean, Std and Median number of size $k = 5$ graphlets per graph.

Dataset	Mean	Std	Median
NCI109	1128404.0613	5330478.99592	65780.0
NCI1	1164713.35912	5896578.49179	80730.0
PTC	732605.578488	3243397.46737	26334.0
MUTAG	16221.1702128	19687.3766842	7378.0

7. DISCUSSION

We presented a novel framework for smoothing normalized graphlet-frequency vectors inspired by smoothing techniques from natural language processing. Although our models are inspired by work done in language models, they are fundamentally different in the way they define a fallback base distribution. We believe that our framework has applicability beyond graph kernels, and can be used in any structural setting where one can naturally define a relationship such as the DAG that we defined in Figure 4. We are currently investigating the applicability of our framework to string kernels [22, 23, 37] and tree kernels [8]. It is also interesting to investigate if our method can be extended to other graph kernels such as random walk kernels. Our framework is also applicable to node-labeled graphs since they also suffer from similar sparsity issues. We leave the application of our framework to labelled graphs to an extended version of this paper. We are also investigating better strategies for tuning the hyper-parameters of the Pitman-Yor kernels.

References

[1] Mathieu Bastian, Sebastien Heymann, and Mathieu Jacomy. Gephi: an open source software for exploring and manipulating networks. In *ICWSM*, pages 361–362, 2009.

[2] D. Bonchev and D. H. Rouvray, editors. *Chemical Graph Theory: Introduction and Fundamentals*, volume 1. Gordon and Breach Science Publishers, London, UK, 1991.

[3] K. M. Borgwardt, C. S. Ong, S. Schönauer, S. V. N. Vishwanathan, A. J. Smola, and H.-P. Kriegel. Protein function prediction via graph kernels. In *Proceedings of Intelligent Systems in Molecular Biology (ISMB)*, Detroit, USA, 2005. `http://www.stat.purdue.edu/~vishy/papers/BorOngSchVisetal05.pdf`.

[4] K. M. Borgwardt, C. S. Ong, S. Schonauer, S. V. N. Vishwanathan, A. J. Smola, and H. P. Kriegel. Protein function prediction via graph kernels. *Bioinformatics (ISMB)*, 21 (Suppl 1):i47–i56, Jun 2005.

[5] Karsten M. Borgwardt and Hans-Peter Kriegel. Shortest-path kernels on graphs. In *Proc. Intl. Conf. Data Mining*, pages 74–81, 2005.

[6] C.C. Chang and C.J. Lin. *LIBSVM: a library for support vector machines*, 2001. Software available at `http://www.csie.ntu.edu.tw/~cjlin/libsvm`.

[7] Stanley F Chen and Joshua Goodman. An empirical study of smoothing techniques for language modeling. In *Proceedings*

of the 34th annual meeting on Association for Computational Linguistics, pages 310–318. Association for Computational Linguistics, 1996.

[8] M. Collins and N. Duffy. Convolution kernels for natural language. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*, pages 625–632, Cambridge, MA, 2001. MIT Press.

[9] A. K. Debnath, R. L. Lopez de Compadre, G. Debnath, A. J. Shusterman, and C. Hansch. Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. correlation with molecular orbital energies and hydrophobicity. *J Med Chem*, 34:786–797, 1991.

[10] P. D. Dobson and A. J. Doig. Distinguishing enzyme structures from non-enzymes without alignments. *J Mol Biol*, 330 (4):771–783, Jul 2003.

[11] T. Gärtner, P.A. Flach, and S. Wrobel. On graph kernels: Hardness results and efficient alternatives. In B. Schölkopf and M. K. Warmuth, editors, *Proc. Annual Conf. Computational Learning Theory*, pages 129–143. Springer, 2003.

[12] Sharon Goldwater, Tom Griffiths, and Mark Johnson. Interpolating between types and tokens by estimating power-law generators. *Advances in neural information processing systems*, 18:459, 2006.

[13] Sharon Goldwater, Thomas L Griffiths, and Mark Johnson. Producing power-law distributions and damping word frequencies with two-stage language models. *The Journal of Machine Learning Research*, 12:2335–2382, 2011.

[14] Aric Hagberg, Pieter Swart, and Daniel S Chult. Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Laboratory (LANL), 2008.

[15] David Haussler. Convolution kernels on discrete structures. Technical Report UCS-CRL-99-10, UC Santa Cruz, 1999.

[16] T. Horvath, T. Gärtner, and S. Wrobel. Cyclic pattern kernels for predictive graph mining. In *Proceedings of the International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 158–167, 2004.

[17] Jaz Kandola, Thore Graepel, and John Shawe-Taylor. Reducing kernel matrix diagonal dominance using semi-definite programming. In *Proc. Annual Conf. Computational Learning Theory*, volume 2777 of *Lecture Notes in Computer Science*, pages 288–302, Washington, DC, 2003.

[18] H. Kashima, K. Tsuda, and A. Inokuchi. Kernels for graphs. In K. Tsuda, B. Schölkopf, and J.P. Vert, editors, *Kernels and Bioinformatics*, pages 155–170, Cambridge, MA, 2004. MIT Press.

[19] J. Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM*, 46(5):604–632, November 1999.

[20] R. Kneser and H. Ney. Improved backing-off for M-gram language modeling. In *Proc. ICASSP '95*, pages 181–184, Detroit, MI, May 1995.

[21] Risi Kondor and Karsten Borgwardt. The skew spectrum of graphs. In *Proceedings of the International Conference on Machine Learning*, pages 496–503. ACM, 2008.

[22] C. Leslie and R. Kuang. Fast kernels for inexact string matching. In *Proc. Annual Conf. Computational Learning Theory*, 2003.

[23] C. Leslie, E. Eskin, and W. S. Noble. The spectrum kernel: A string kernel for SVM protein classification. In *Proceedings of the Pacific Symposium on Biocomputing*, pages 564–575, Singapore, 2002. World Scientific Publishing.

[24] C. D. Manning, P. Raghavan, and H. Schütze. *Introduction to Information Retrieval*. Cambridge University Press, 2008. URL `http://nlp.stanford.edu/IR-book/`.

[25] Brendan D McKay. Nauty userâĂŹs guide (version 2.4). *Computer Science Dept., Australian National University*, 2007.

[26] L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford Digital Library Technologies Project, Stanford University, Stanford, CA, USA, November 1998.

[27] J. Pitman and M. Yor. The two-parameter poisson-dirichlet distribution derived from a stable subordinator. *Annals of Probability*, 25(2):855–900, 1997.

[28] N. Przulj. Biological network comparison using graphlet degree distribution. In *2006 European Conference on Computational Biology (ECCB)*, September 2006.

[29] J. Ramon and T. Gärtner. Expressivity versus efficiency of graph kernels. Technical report, First International Workshop on Mining Graphs, Trees and Sequences (held with ECML/PKDD'03), 2003.

[30] R. Sharan and T. Ideker. Modeling cellular machinery through biological network comparison. *Nature Biotechnology*, 24(4): 427–433, Apr 2006.

[31] Nino Shervashidze and Karsten Borgwardt. Fast subtree kernels on graphs. In *Neural Information Processing Systems*, 2010. URL `http://books.nips.cc/papers/files/nips22/NIPS2009_0533.pdf`.

[32] Nino Shervashidze, S. V. N. Vishwanathan, Tobias Petri, Kurt Mehlhorn, and Karsten Borgwardt. Efficient graphlet kernels for large graph comparison. In Max Welling and David van Dyk, editors, *Proc. Intl. Conference on Artificial Intelligence and Statistics*. Society for Artificial Intelligence and Statistics, 2009.

[33] Nino Shervashidze, Pascal Schweitzer, Erik Jan Van Leeuwen, Kurt Mehlhorn, and Karsten M Borgwardt. Weisfeiler-lehman graph kernels. *The Journal of Machine Learning Research*, 12:2539–2561, 2011.

[34] Q. Shi, J. Petterson, G. Dror, J. Langford, A. J. Smola, A. Strehl, and S. V. N. Vishwanathan. Hash kernels. In M. Welling and D. van Dyk, editors, *Proc. Intl. Conference on Artificial Intelligence and Statistics*. Society for Artificial Intelligence and Statistics, 2009.

[35] Yee Whye Teh. A hierarchical bayesian language model based on pitman-yor processes. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*, pages 985–992. Association for Computational Linguistics, 2006.

[36] H. Toivonen, A. Srinivasan, R. D. King, S. Kramer, and C. Helma. Statistical evaluation of the predictive toxicology challenge 2000-2001. *Bioinformatics*, 19(10):1183–1193, July 2003.

[37] S. V. N. Vishwanathan and A. J. Smola. Fast kernels for string and tree matching. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 569–576. MIT Press, Cambridge, MA, 2003.

[38] S. V. N. Vishwanathan, Karsten Borgwardt, and Nicol N. Schraudolph. Fast computation of graph kernels. In B. Schölkopf, J. Platt, and T. Hofmann, editors, *Advances in Neural Information Processing Systems 19*, Cambridge MA, 2007. MIT Press.

[39] S. V. N. Vishwanathan, Nicol N. Schraudolph, Imre Risi Kondor, and Karsten M. Borgwardt. Graph kernels. *Journal of Machine Learning Research*, 2010. URL `http://www.stat.purdue.edu/~vishy/papers/VisSchKonBor10.pdf`. In press.

[40] Nikil Wale, Ian A Watson, and George Karypis. Comparison of descriptor spaces for chemical compound retrieval and classification. *Knowledge and Information Systems*, 14(3): 347–375, 2008.

[41] Takashi Washio and Hiroshi Motoda. State of the art of graph-based data mining. *SIGKDD Explorations*, 5(1):59–68, 2003.

[42] C. Zhai and J. Lafferty. A study of smoothing methods for language models applied to information retrieval. *ACM Trans. Inf. Syst.*, 22(2):179–214, 2004. ISSN 1046-8188. doi: http://doi.acm.org/10.1145/984321.984322.